

REMARKS

1. Summary of Office Action

In the Office Action dated November 5, 2004, the Examiner stated that claims 1-8 and 10-29 were pending. Claim 9 was previously canceled, and claims 1, 22, and 25 were previously amended. Claims 1, 22, and 25 have currently been amended. No new matter has been added. Currently pending are claims 1-8 and 10-29.

All currently pending claims were rejected by the Examiner. Of the claims the Examiner (1) rejected claims 22-29 under 35 U.S.C. § 102(e) as being anticipated by Mulligan, U.S. Patent 6,212,190 (hereinafter, "Mulligan"); and (2) rejected claims 1-8 and 10-21 under 35 U.S.C. § 103(a) as being unpatentable over Mulligan in view of Matsuzono, U.S. Patent 5,809,254 (hereinafter, "Matsuzono").

After careful review of the pending claims and the cited references, Applicants submit all pending claims as currently in condition for allowance and respectfully request favorable reconsideration in view of the following remarks.

2. Response to 35 U.S.C. § 102(e) Rejection

Claims 22-29 were rejected under 35 U.S.C. § 102(e) as being anticipated by Mulligan. Of these claims, dependent claims 23 and 24 derive from independent claim 22, and dependent claims 26-29 derive from independent claim 25. Claims 22 and 25 have been amended to better claim the invention.

Amended claim 22 is directed toward a method for reducing message fragmentation between a data source and a data receiver connected by a network. This method comprises the steps of

intercepting a communication, predicting a maximum segment size (MSS), sending a signal to the data source and the data receiver, storing a determined MSS, and inserting the determined MSS into subsequent connection announcements. In the step of sending a signal to the data source and the data receiver, Applicants specifically state the use of a signal “with a no-fragment option set” (claim 22). In the specific embodiment claimed, it is necessary that the no-fragment option be set for the signal so that the network device attempting to reduce message fragmentation may receive network information regarding the appropriate maximum transfer unit (MTU) of the connection or connections between the data source and data receiver (specification, page 21, lines 1-15). If the no-fragment option is not set, the signal may be fragmented by network devices and may successfully reach the data source and data receiver regardless of its size; in this event, no useful network information pertaining to the MTU of the network links is obtained. Additionally, the method of claim 22 states the step of sending the signal with the no-fragment option set to both the data source and the data receiver. This indicates the ability of the method to be used in conjunction with a network device that bridges two different networks, with the data source located on one network and the data receiver being located on the other. The method of claim 22 also states the insertion of the determined MSS into communications, and may be accomplished by such means as inserting a value for the MSS into the options field of a TCP header (specification, page 16, lines 15-20), or by modifying the size of the announcement message (specification, page 9, lines 18-22). As a result, the method of claim 22 functions by adjusting the initial messages used to establish a connection between the data source and the data receiver.

Mulligan does not teach a method of reducing message fragmentation between a data source and a data receiver connected by a network consistent with the method of claim 22. Specifically, Mulligan does not teach the method step of sending a signal to both the data source and the data

receiver, and does not teach the step of adjusting the MSS of the initial message used to setup the connection between the source host and the destination host. Rather, Mulligan determines the MTU of a link in a single direction (i.e. from the source to the destination) using a technique called "Path MTU Discovery" (col. 3, lines 45-65; col. 8, lines 46-57). As described by Mulligan, although this technique utilizes datagram packets with the don't fragment (DF) bit set, the datagram packets are sent directly from the source host to the destination host, and are not sent to both the source host and the destination host. In addition, this process does not require the step of intercepting a connection announcement between the source and destination hosts, which is explicitly required by the method of claim 22. As a result, Mulligan does not teach a method that is capable of bridging two networks with different MTUs, which is a unique aspect of the current invention. In addition, no methods taught by Mulligan concern connection announcements between nodes in any capacity. Rather, the methods described by Mulligan work on the assumption that MTUs for network segments can initially be determined (Mulligan, column 8, lines 34-35), and then address the problem of message fragmentation by determining the best size to use for fragmenting messages that are greater than a given MTU (Mulligan, column 9, lines 38-48; column 10, lines 32-41). In contrast, the method of claim 22 addresses this problem by ensuring that the messages generated by the source and destination hosts are always less than or equal to the MTU of the network, thereby preventing any message fragmentation. In short, Mulligan seeks to optimize message fragmentation while the current invention seeks to avoid fragmentation entirely. Therefore, Mulligan does not teach the method of claim 22.

Claim 25 of the current application is directed towards an apparatus for reducing message fragmentation between a data source and a data receiver. In a similar manner that claim 22 seeks to address the issue by reducing message fragmentation, the apparatus of claim 25 modifies

communications between a source and host so that an MSS is utilized by the source and destination hosts that permits transmission through the network with little or no fragmentation. The apparatus functions by intercepting communications between the source and destination hosts, adjusting the MSS indicated by the communications, and then sending the modified communications to the source and destination hosts. With the modified MSS communications, the source and destination hosts can adjust their message sizes to the STE-provided MSS, thereby avoiding fragmentation.

As discussed above, Mulligan describes a network device that determines if the size of a message is larger than a determined MTU value. If so, the device fragments the message according to an optimization scheme (Mulligan, column 9, lines 38-66). However, unlike the current invention, the network device takes no action to correct this issue from occurring in the future by ensuring that subsequent messages generated by the source or destination host are already less than or equal to the size of the determined MTU value. As a result, Mulligan does not teach the apparatus of claim 25.

Applicants submit that due to the substantial differences between the methods and devices taught by Mulligan and those of independent claims 22 and 25, evidence of novelty for these claims is shown in this case. Furthermore, Applicants submit that dependent claims 23 and 24, which derive from claim 22, and dependent claims 26-29, which derive from claim 25, are also shown to be novel.

In light of the above remarks, Applicant respectfully requests withdrawal of the 35 U.S.C. § 102(e) rejections.

3. Response to 35 U.S.C. § 103(a) Rejection

Claims 1-8 and 10-21 were rejected under 35 U.S.C. § 102(a) and (e) as being anticipated by Smith. Of these claims, dependent claims 2-8 and 10 derive from independent claim 1, and

dependent claims 12-21 derive from independent claim 11. Claim 1 has been amended to better claim the invention.

Amended claim 1 is directed toward a method for managing message size in a network connection between a data source and data receiver, wherein the method comprises the steps of receiving an announcement establishing a connection between a data source and a data receiver, and then changing the maximum segment size in the announcement to a determined maximum segment size. Specifically, this method does not involve the fragmentation of datagrams in order to alleviate problems associated with message fragmentation.

Mulligan does not teach a method where connection announcements are modified by changing the maximum segment size denoted in the announcement. As described above, the invention of Mulligan is directed towards mitigation of fragmentation effects by finding an optimal size for fragmenting messages. Of the specific references to Mulligan cited by the Examiner in this regard, none detail the modification of the MSS in the manner suggested by the method of claim 1: column 5, lines 27-34 describe a simple network setup; column 7, lines 55-67 provide general information indicating that messages larger than an MTU need to be fragmented or dropped; and column 9, lines 38-48 describe the determination of the size of the datagram and the use of optimized message fragmentation. As a result, Mulligan does not describe the step of changing the MSS in the announcement message to a determined MSS. Therefore Mulligan does not teach all of the steps of the method of claim 1.

Matsuzono fails to cover this deficiency. Matsuzono describes a system that resides on a host computer, wherein the system accepts MSS requests and issues MSS queries (Matsuzono, column 5, lines 3-31). However, the system does not intercept communications between source and destination hosts. In addition, the system taught by Matsuzono does not describe the step of

changing the MSS in the announcement message to a determined MSS. Because neither Mulligan nor Matusuzono teaches the step of changing the maximum segment size in the announcement of the first connection to a determined maximum segment size, all aspects of the current invention are not taught by Mulligan in view of Matsuzono. As a result, Mulligan in view of Matsuzono does not establish a *prima facie* case of obviousness with respect to claim 1.

Claim 11 is directed toward a method for reducing message fragmentation for a connection between a data source and data receiver, wherein the method comprises the steps of receiving a first message fragment of a first connection between the data source and the data receiver, storing the MSS size of the first message, resetting the first connection and initiating the creation of a second connection, and the placing the MSS into the announcement of the second connection.

Mulligan does not teach a method involving the step of resetting a connection between the data source and the data receiver. As described above, the invention of Mulligan is directed towards mitigation of fragmentation effects by finding an optimal size for fragmenting messages. Of the specific references to Mulligan cited by the Examiner in this regard, none detail the modification of the MSS in the manner suggested by the method of claim 11: column 5, lines 27-34 describe a simple network setup; column 7, lines 55-67 provide general information indicating that messages larger than an MTU need to be fragmented or dropped; and column 9, lines 38-48 describe the determination of the size of the datagram and the use of optimized message fragmentation. No indication of resetting a connection in order to establish a second connection is described, either explicitly or implicitly, by Mulligan. Additionally, Mulligan does not teach the step of placing a MSS size into an announcement of a connection. Of the specific references to Mulligan cited by the Examiner in this regard, none detail the insertion of the MSS into a connection announcement in the manner suggested by claim 11: column 6, lines 18-29 describe the ability of a processor to send

and receive data; column 7, lines 55-67 provide general information indicating that messages larger than an MTU need to be fragmented; and column 9, lines 7-21 describe a process to determine an optimal MTU. As a result, Mulligan does not describe either the step of intentionally resetting a connection to initiate a second connection, or placing the MSS into the announcement of the second connection. Therefore Mulligan does not teach all of the steps of the method of claim 11.

Matsuzono fails to cover these deficiencies. Matsuzono describes a system that resides on a host computer, wherein the system accepts MSS requests and issues MSS queries (Matsuzono, column 5, lines 3-31). However, the system does not reset connections between source and destination nodes. In addition, the system taught by Matsuzono does not describe the step of placing a determined MSS in a connection announcement message. Because neither Mulligan nor Matsuzono teaches the steps of intentionally resetting a connection to initiate a second connection, or placing the MSS into the announcement of the second connection, all aspects of the current invention are not taught by Mulligan in view of Matsuzono. As a result, Mulligan in view of Matsuzono does not establish a *prima facie* case of obviousness with respect to claim 11.

Applicants submit that due to the substantial differences between the methods and devices taught by Mulligan in view of Matsuzono and those of independent claims 1 and 11, it is sufficiently shown that these claims are novel and non-obvious. Furthermore, Applicants submit that dependent claims 2-8 and 10, which derive from claim 1, and dependent claims 12-20, which derive from claim 11, are also shown to be novel and non-obvious.

In light of the above remarks, Applicant respectfully requests withdrawal of the 35 U.S.C. § 103(a) rejections.

CONCLUSION

In light of the above amendments and remarks, Applicants submit that the present application is in condition for allowance and respectfully requests notice to that effect. The Examiner is requested to contact Applicants' representative below at (312) 913-3303 if any questions arise or he may be of assistance to the Examiner.

Date: 8/17/05

Respectfully Submitted,

By:



Brian R. Harris

Reg. No. 45,900

McDonnell, Boehnen

Hulbert & Berghoff LLP

300 South Wacker Drive

Chicago, IL 60606

Phone: (312) 913-0001

Fax: (312) 913-0002